

REMARKS

This Preliminary Amendment is filed in response to the Final Office Action mailed on July 11, 2006, and is filed herewith a Request for Continuing Examination. All objections and rejections are respectfully traversed.

Claims 1-22, 38, 40-41, 43-44, 46, 55-81 are currently pending.

Claim Rejections – 35 USC § 102

At paragraphs 4-5 of the Office Action, claims 1, 3, 8-12, 14-22, 55-58, and 61-81 are rejected under 35 U.S.C. §102 as being anticipated by Stallmo et al., US Patent No. 6,052,759, issued on April 18, 2000, hereinafter Stallmo.

The present invention, as set forth in representative claim 1, comprises in part:

1. A method for enabling parity declustering in a balanced parity array of a storage system, where an operating system performs the method comprising the steps of:

combining a plurality of unbalanced parity stripe arrays, each unbalanced parity stripe array storing an unequal number of parity blocks per disk, to form the balanced parity array, the balanced parity array storing substantially the same number of parity blocks on all disks, each unbalanced parity stripe array and the balanced parity array having parity blocks on a set of storage devices that are disjoint from a set of storage devices storing data blocks; and
distributing assignment of storage devices to parity groups throughout the balanced parity array.

By way of background, Stallmo describes a redundant array of independent disks (RAID) system for organizing data across the array when disks are added and removed. Stallmo's process of combining parity (Fig. 12, 14 and 15) onto a single disk requires

swapping parity blocks and data blocks to go from a RAID 5 implementation to a RAID 3 implementation. (Col. 14, lines 19-22). Stallmo performs this change to allow for a disk to be added or deleted from the array. Furthermore, Stallmo organizes the disks into “rectangles” where each rectangle has the same number of blocks on each disk with the first “rectangle” equal to the number of blocks in the disk with the smallest number of blocks. The disk may have more than one rectangle with except for the smallest disk. (Col. 8, lines 45-51). The “rectangles” are organized into “squares”, where the number of blocks in each “square” is equal to the number of disks in the rectangle multiplied by the depth. (Col. 9, lines 14-20) The depth is set by the sub-system software and equal to the number of blocks that are necessary for each read or write operation. (Col. 2, lines 13-17). Stallmo organizes the data into rectangles to transform unstriped unprotected data into striped data or if data and parity are involved to transform from protected unstriped to protected striped.

Applicant respectfully urges that Stallmo does not describe Applicant’s novel *combining a plurality of unbalanced parity stripe arrays, each unbalanced parity stripe array storing an unequal number of parity blocks per disk, to form the balanced parity array, the balanced parity array storing substantially the same number of parity blocks on all disks, each unbalanced parity stripe array and the balanced parity array having parity blocks on a set of storage devices that are disjoint from a set of storage devices storing data blocks and distributing assignment of storage devices to parity groups throughout the balanced parity array.* In further detail, Applicant’s invention is combining stripes of data and parity to form a super stripe. The storage devices that contain data are separate (*disjoint*) from the storage devices that contain parity. The plurality of stor-

age devices containing parity form a parity array, which is a subset of all the disks in the RAID array. Specifically, the system combines an *unbalanced stripe* which has *an unequal number of parity blocks per disk* with at least one other unbalanced stripe having *an unequal number of parity blocks per disk* to form a balanced parity array of which stores *substantially the same number of parity blocks on all disks*.

By *distributing assignment of storage devices to parity groups throughout the balanced array*, enables Applicant's invention to handle single and double disk failures. Additionally, the number of accesses to each number of disks is reduced by distributing the parity for two separate sets of data across all disks. (See Specification page 28, line 14 to page 29, line 2, Fig. 14).

In contrast, Stallmo does not describe use of a *parity array*. Stallmo discloses either spreading parity across all disks of the RAID array (i.e. RAID 5) or a single disk containing all parity (i.e. RAID 3). When Stallmo discloses combining parity into a single disk for adding or removing a disk (switching from RAID 5 to RAID 3), Stallmo describe swapping parity disk assignments with data assignments. In other words, Stallmo is moving data to align all parity blocks on a single disk. There is no disclosure of using multiple parity disks in Stallmo, nor any suggestion because a single disk of parity is sufficient for single disk failures. Stallmo does not suggest or disclose double disk failures.

Furthermore, the Examiner states a "balanced array" is shown by Stallmo, Col. 8, lines 45-51, which states:

“One goal of the system is to allow disks of varying sizes, that is having varying numbers of data blocks, to be managed and to assure that all the blocks on each disk are available to the host computer. When multiple disks are managed, they are organized into multiple "rectangles", where each rectangle has a set of disks that all contain the same number of blocks. The number of rectangles needed is determined by the number of disks that have varying sizes.”

In reference to the statement above, Stallmo describes a system of organizing a plurality of blocks on all disks into rectangles, where the disks are unstriped. The first rectangle is merely equal to the size of the disk with the lowest number of blocks. The second rectangle is merely equal to the difference between the disk with the next lowest number of blocks and the disk with the lowest number of blocks, and so until you reach the disk with the largest number of blocks. The rectangles are then converted into stripes based on the number of disks in each rectangle to be used in either (RAID 0, 1, 3, or 5). However, there is no combination of unbalanced arrays in Stallmo because each of the rectangles have the same number of blocks per disk, and the system steps through each rectangle separately. (See Stallmo Fig. 12).

The blocks in Stallmo can be data or parity. There is no disclosure of only organizing parity blocks into balanced arrays. Additionally, Stallmo does not disclose or suggest use of a plurality of disks for storing parity because Stallmo only teaches of single disk failure. Additionally, Stallmo teaches use of spreading the parity across all disks (i.e. RAID 5), which teaches away from Applicant's parity blocks stored *on a set of storage devices that are disjoint from a set of storage devices storing data blocks*.

Additionally, the Examiner further states that Stallmo shows organizing a balanced array in Col. 9, lines 15-67, which states:

“To accomplish these goals, each rectangle is divided into a set of "squares". A square is a portion of the set of disks contained within a rectangle. The number of blocks in each square is equal to the number of disks in the rectangle multiplied by the depth being used by the rectangle. Each square typically starts at the same logical block number on each disk.

Since the number of blocks in a rectangle is not necessarily an even multiple of the number of blocks in a square, there may be a "partial" at the end of the rectangle, and this partial portion contains the remaining blocks in each disk that cannot fit in a square. These partial blocks do not participate in the striping operation, described below, and thus remain un-striped. They will have data protection, however, since parity can be maintained with an un-striped configuration.

FIG. 5 shows an example of a rectangle and some squares that fit into the rectangle. Referring to FIG. 5, four disks 502, 504, 506, and 508 are shown, wherein the disks comprise a rectangle containing 1000 blocks on each disk. In this example, the depth is four blocks, and since there are four disks, each square contains 16 blocks on each disk. In this example, the rectangle contains 62 squares, and there are eight blocks left over on each disk. These left over blocks comprise a partial.

The squares organization is used to allow data to be striped and un-striped across disks. Since each square has the same number of rows and columns, wherein one row is a depth's worth of blocks, and there is one column per disk, matrix transposition is used on a square to stripe and unstripe data blocks, as will be described below with respect to FIGS. 11-16.

The management of data on the disks of the array is layered. At the first level is the management of striping of data blocks and possibly parity blocks. The first level of management is also responsible for sparing and reconstruction operations. This level is called transparent RAID. The second level of management is adaptive RAID, as will be described below.

In transparent RAID, the only configuration information the host/user/system administrator can specify is that an added disk is to be used as a data disk, a parity disk or a spare disk. The disk array uses a disk as a data disk if the type of disk is not defined. The host/user/system administrator can also specify the depth, that is the number of blocks written on a specific disk before writing moves to the next disk.

In transparent RAID, the data blocks on each disk and the parity blocks, if a parity disk is being used, are automatically striped across all of

the managed disks in the set. When a new disk is added to an existing set of managed disks, all the data on the existing disks is re-stripped across all the disks including the new disk. The blocks on the new disk are also striped across all of the disks in the managed set.” (emphasis added)

In reference to the above statement, Stallmo describes organizing most of the unstriped data blocks in each rectangle into stripes. The remaining blocks that do not fit into any of the squares of each rectangle are not striped. They are part of an incomplete or unbalanced stripe of data or data and parity. Stallmo only can create balanced stripes (stripes with the same number of blocks per disk) if there is a specific multiple of blocks per rectangle based on the size of the depth preset. Additionally, these balanced arrays of are mostly data, with the parity rotatably striped throughout all the disks of the RAID array or all the parity is stored on a single disk. There is no disclosure in Stallmo of balancing parity blocks on a plurality of storage devices where each disk stores only parity because Stallmo only teaches of at most one disk of parity or rotating the parity through all the disks of the array.

Accordingly, Applicant respectfully urges that Stallmo is legally insufficient to anticipate the present claims under 35 U.S.C. §102 because of the absence of the Applicant’s claimed novel *combining a plurality of unbalanced parity stripe arrays, each unbalanced parity stripe array storing an unequal number of parity blocks per disk, to form the balanced parity array, the balanced parity array storing substantially the same number of parity blocks on all disks, each unbalanced parity stripe array and the balanced parity array having parity blocks on a set of storage devices that are disjoint*

from a set of storage devices storing data blocks and distributing assignment of storage devices to parity groups throughout the balanced parity array.

Claim Rejections – 35 USC § 103

At paragraph 4 of the office Action, claims 38, 40-41, 43-44, and 46 were rejected under 35 U.S.C. §102 as being anticipated by Burton et al. US Patent Application Publication 2003/0074527, filed on April 17, 2003, in view of Stallmo.

Applicant respectfully points out that §102 rejection is improper because is based on a combination of two references. Applicant will assume the rejection is a §103 rejection, and respond accordingly.

The present invention, as set forth in representative claim 38, comprises in part:

38. A method for declustering a parity array having a plurality of storage devices, where an operating system performs the method comprising the steps of:

assigning a first plurality of data and parity blocks to a first group, where the data blocks are stored on a plurality of storage devices that are disjoint from a plurality of storage devices storing parity blocks to form a first parity group, wherein the parity blocks are striped across the plurality of storage devices in the first parity group with an uneven number of parity blocks per each storage device;

assigning a second plurality of data and parity blocks to a second group, where the data blocks are stored on a plurality of storage devices that are disjoint from a plurality of storage devices storing parity blocks to form a second parity group, wherein the parity blocks are striped across the plurality of storage devices in the second parity group with an uneven number of parity blocks per each storage device, wherein the first and second parity groups being independent from each other; and

combining the first parity group and the second parity group to form a balanced parity array, the balanced parity array storing substantially the same number of parity blocks on the storage devices configured to store parity.

By way of background, Burton describes a system for a user to create a span of disks. The system wants a balanced number of disks in each span to maximize cache utilization. With more disks per span, fewer stripes are necessary to store all the cache information.

Applicant respectfully urges that Burton and Stallmo taken alone or in combination do not teach or suggest Applicant's claimed novel step of *combining the first parity group and the second parity group to form a balanced parity array, the balanced parity array storing substantially the same number of parity blocks on the storage devices configured to store parity*. In further detail, Applicant's invention arranges the disks to have data disks and parity disks, where the parity disks form a parity array. A first parity group is stored on the parity disks of the parity array. The second parity group is also stored on the parity disks of the parity array. There are gaps in locations of where parity is stored for example with an even number of disks (n) there can be no delta of $n/2$ because if two disks were half the length of the array apart were lost then, the data blocks members of the same two parity sets are lost, thereby obviating recovery of the lost data. (Specification page 20, line 1-26). When the two groups are combined, a balanced parity array forms combining an unbalanced parity array from first parity group and the second parity group.

In contrast, there is no disclosure in either Stallmo or Burton of using a parity array, or a plurality of disks for storing parity. Specifically, there is no suggestion in Stallmo because it only discloses single disk failure which only requires a single disk of

parity to protect the data. Additionally, Stallmo also discloses rotating the parity through all disks instead of just storing all the parity on a single disk. Furthermore, Burton only describes balancing the number of disks in each span and not combining parity groups to form a balanced array with approximately the same number of blocks per all disks.

Accordingly, Applicant respectfully urges that the Stallmo patent and the Burton patent, either taken singly or taken in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103 because of the absence in each of the cited patents of Applicant's claimed novel *combining the first parity group and the second parity group to form a balanced parity array, the balanced parity array storing substantially the same number of parity blocks on the storage devices configured to store parity.*

At paragraph 5 of the Office Action, claims 2, and 4-6 were rejected under 35 U.S.C. §103 as being unpatentable over Stallmo, in view of Baylor et al., US Patent No. 5,862,158, hereinafter Baylor.

At paragraph 6 of the Office Action, claims 7, 13, and 59-60 were rejected under 35 U.S.C. §103 as being unpatentable over Stallmo, in view of Baylor, and in further view of Karr, US Patent No. 3,993,862, hereinafter Karr.

Applicant respectfully notes that claims 2, 4-7, 13, and 59-60 are dependent claims that depend from independent claims believed to be in condition for allowance.

Accordingly, claims 2, 4-7, 13, and 59-60 are also believed to be in condition for allowance.

In the event that the Examiner deems personal contact desirable in the disposition of this case, the Examiner is encouraged to call the undersigned attorney at (617) 951-3067.

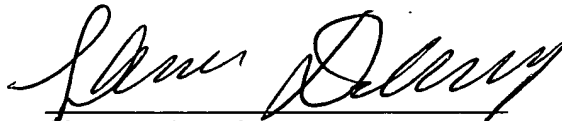
All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims.

The Applicant respectfully solicits favorable action.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,



Shannen C. Delaney
Reg. No. 51,605
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500